

TEXAS MEDICAL CENTER NASA/JOHNSON SPACE CENTER
COOPERATIVE AGREEMENT PROGRAM NCC 9-36, ROUND 1

COVER SHEET FOR FINAL REPORT

Please follow this format:

Name of Subcontractor: Marc J. Fenstermacher, M. D.

Title: Assistant Professor

Institution: University of Texas M. D. Anderson Cancer Center

Name of Project: An Efficient, Lossless Database for Storing and
Transmitting Medical Images

Amount of Grant: ~~899,000~~

* Amount Spent, if Different from Amount Granted: ~~899,000~~

Date Project Was Completed: June 30, 1998

A. Mastromarino
Grants Officer: Anthony J. Mastromarino, Ph. D.

Title: Associate Vice President for Faculty Affairs

Phone: 713-792-3391

Fax: 713-745-4277

1. INTRODUCTION

This research aimed in creating new compression methods based on the central idea of *Set Redundancy Compression* (SRC). Set Redundancy refers to the common information that exists in a set of similar images. SRC compression methods take advantage of this common information and can achieve improved compression of similar images by reducing their Set Redundancy. The current research resulted in the development of three new *lossless* SRC compression methods: MARS (Median-Aided Region Sorting), MAZE (Max-Aided Zero Elimination) and MaxGBA (Max-Guided Bit Allocation).

2. THE TEST IMAGES

Different image types were considered for this research. The evaluation criteria included image availability, need for improved compression, maximization of end-user benefits, and fitness of images to the SRC scheme. After considering several modalities and image types and running preliminary tests, it was concluded to use chest X-rays, brain CT, and brain MR images. A test image database was created by retrieving 100 chest X-rays, 51 brain CT, and 57 brain MR images from M.D. Anderson's image archives. This database was further expanded by adding smoothed versions of the brain CT images, and larger size versions of the chest X-ray images.

3. IMAGE REGISTRATION

Most of the SRC methods perform better when the images are registered in a standard position, orientation, and size. Note that we can maintain the lossless property of compression by performing “inverted” registration (i.e. instead of registering the original images on a given statistical image, we keep the original images intact and instead register the statistical image on each original image before compression). The statistical images that SRC methods use are the “median” image (in the MARS method) and the “maximum” image (in the MAZE and MaxGBA methods). The median image is formed by using for each pixel position the median value from the same pixel position across all given images. Similarly, the maximum image is formed by using the maximum value for each pixel across all images. In our tests we used registration for the MR and CT images, but not on the X-ray images.

4. IMAGE CLUSTERING

The SRC compression methods work with groups of highly similar images. In a given image database we need to form these groups before we apply any SRC method. Forming these similar image groups corresponds to the image clustering problem. We developed an automated process based on *Genetic Algorithms* to perform image clustering in our test image database. The resulting image groups were used for all subsequent tests.

5. TESTS

A total of 341 tests were run (11 image types x 31 methods). The different image types were CT, MR, X-rays, smoothed images, differential images, images with varying sizes, etc. For each image type, we run tests with 31 different methods that include all of our newly developed methods combined with either Lempel-Ziv (compress), LZW (gzip) or Arithmetic (bit-based, char-based, or word-based) as the entropy-encoder, and also standard methods used for comparison purposes (bit- char- and word-based Arithmetic compression, Lempel-Ziv, LZW, and Calic).

6. THE SOFTWARE

Software was written to extract image data from files stored in MEDVISION and DICOM formats. These two image file formats are widely used in MDACC. Code was developed to remove patient data from the test images in order to preserve the patient confidentiality throughout our tests. Unix scripts were written to facilitate running tests in batch mode. Software was also developed to perform basic image processing operations in order to assist testing and exploration of image properties. A Genetic Algorithms program was written to perform image clustering. Finally, software was developed to implement the MARS, MAZE and MaxGBA methods. All software was designed using the object-oriented design methodology and was implemented in C++ on a Unix platform.

7. THE NEW SRC COMPRESSION METHODS

Three new lossless Set Redundancy Compression methods were developed in this research: the MARS (Median-Aided Region Sorting) method, the MAZE (Max-Aided Zero Elimination) method and the MaxGBA (Max-Guided Bit Allocation) method.

7.1 *The MARS method*

The Median-Aided Region Sorting (MARS) method starts with converting the input images into their “differential” versions replacing pixel values with their differences from a predicted value (the prediction is based on a scheme similar to the one used in the Calic method). The “median” image is then created from these differential images, and it is segmented into regions of same value. These regions are subsequently used as a guide to “sort” the pixel values in each input image: based on the “region map” from the median image, we rearrange the pixels in each input image so that pixels from the value-0 regions are stored first, followed by the pixels from the value-1 regions, and so on. The “sorting” would be perfect only if it was applied to the median image itself; in all other images it is expected to be only approximate, since in general the region map of the median image only approximately represents the real region map of any of the input images. However, although this pixel sorting is not perfect, it can help to improve compression by using any standard entropy-encoder as the last step. The theoretical “best-case” when sorting is perfect was studied experimentally and it was found that it results in almost 90:1 lossless compression. This of course cannot be achieved in practice, however it indicates the potential of the MARS method. The MARS method is lossless.

7.2 *The MAZE method*

Similar to the MARS method, the Median-Aided Zero Elimination (MAZE) method starts with converting the input images into their “differential” versions replacing pixel values with their differences from the predicted values. The “maximum” image is then created from these differential images, and it is used to identify the pixels with 0 value from the input images. Note that if a pixel in the max image has 0 value, then it is *guaranteed* that the same pixel in all input images has also 0 value. These pixels can then safely be removed from all input images, since the “knowledge” to put them back has been captured into the max image. Thus the MAZE method reduces the number of pixels with 0 value from the input images. Since these image are actually differential images, it is expected that they will have a large number of pixels with 0 value, thus the potential size reduction can be significant. The MAZE method is lossless.

7.3 *The MaxGBA method*

The MaxGBA (Maximum-Guided Bit Allocation) method uses the “maximum” image from the set of similar images to guide the bit allocation when encoding pixels of a given input image. The maximum image is created to store the statistical maximum values from a given image set; every pixel of the maximum image contains the largest value from all pixels at the same position across all images in the set.

Therefore, given the maximum image, we know what is the maximum value that each pixel may have (within the given set of images). This information allows us to allocate for every pixel just enough bits to store a value up to the known max value for that pixel. For example, if for some pixel position the maximum image value is 12, then we know that for every image in the set 4 bits will be sufficient to store the value in that pixel position. Note that using a scheme similar to the MAZE method, we can skip pixels for which the corresponding maximum image values are 0. Overall, the smaller values the maximum image has, the greater the savings from using MaxGBA. In order to decrease the values in the maximum image, the input images are pre-processed with a Calic-like prediction scheme to losslessly decrease pixel values; then the maximum image is created, and the MaxGBA method is used, followed by a standard entropy-compression method. Encoded images can be restored with the inverse procedure, using again the maximum image values. The compression that is achieved by MaxGBA is lossless.

In our tests we found that the MaxGBA method almost always performs better than MAZE or MARS. All three methods perform consistently better than LZW, Lempel Ziv, or Arithmetic compression, while also approaching, and in some case exceeding, the performance of CALIC, which many researchers consider to be the best lossless image compression method that exists today. The following table summarizes the results and compares the new SRC methods developed in this research project with Calic and LZW:

	MaxGBA	MAZE	MARS	CALIC	LZW
brain MR	1.897:1	1.830:1	1.829:1	2.014:1	1.382:1
chest X-rays	3.080:1	2.904:1	2.903:1	3.020:1	1.798:1
large chest X-rays	2.894:1	2.799:1	2.799:1	2.836:1	1.776:1
brain CT	4.322:1	4.270:1	4.010:1	4.367:1	2.542:1
smoothed brain CT	5.331:1	5.612:1	5.102:1	4.925:1	2.896:1

8. OTHER RESEARCH DIRECTIONS

8.1 “Factoring out” common pixels

We tested the idea of “factoring out” common pixel values from a set of images. However, the results were not encouraging, since it was found that in practice there are almost no common pixel values (at identical positions), even among only 10 images with high conceptual similarity. It was concluded that it is difficult to exploit conceptual similarity directly into the pixel level, and therefore, we shifted our focus more to methods like MARS or MAZE that attempt to exploit similarities into a “regional” level.

8.2 KLT

Research on KLT (Karhunen-Loeve Transform) was also performed. KLT is theoretically optimal for decorrelating images. We investigated the use of ideas borrowed from KLT to improve the SRC methods. The final conclusion was that KLT would not be useful in our research even though in theory it appeared as an attractive method for compressing similar images. The main reasons for this conclusion are that (a) KLT is mostly powerful when it is used for lossy compression, but in our research we were interested only in lossless methods, (b) KLT was proven to be extremely computationally expensive for our purposes, and (c) the KLT implementations that were not computationally intensive were impractical since they needed to use all images in the set during compressing or decompressing any single image.

8.3 Alternative ways to store pixel differences

We also investigated alternative ways to store pixel differences from predicted values, to improve overall compression (storing pixel differences is part of all SRC compression methods). The most promising idea was to store separately the absolute values and the signs of the differences. For a typical X-ray image, we found that storing absolute values only (without sign) saves about 21KB per image file. However, the file with the signs would be about 29KB (uncompressed), and we found that this file is quite random, therefore it cannot be effectively compressed to produce any meaningful total savings from this technique. The situation is similar for other types of images. Also in all cases this idea would require storing two files per image (the compressed image itself, plus the sign file) which is undesirable. Therefore, we decided to store pixel differences in a plain way, without any further post-processing.

8.4 “Block-matching”

We worked on improving an older SRC method, the CENTROID, by implementing “block-matching”. The idea is to divide each input image into 8x8 square “blocks” of pixels, and then try to find a position on the average image that matches best each block. The CENTROID method can then be used with each block being optimally positioned (smallest pixel differences) on the average image. We developed software to perform block-matching and we explored this idea by running preliminary tests. It was found that the pixel differences in the image blocks were reduced with this method, however, the overall compression wasn’t improved, because of discontinuities in the block boundaries. More research is needed to explore ways to handle these boundary discontinuities.

8.5 “Minimum string edit distance”

On the theoretical side, we started investigating the theory of “minimum string edit distance” and its applications in compressing similar images. Every image can be considered as a “string” of values; furthermore, every string can be transformed into another string by performing operations like deletions, insertions, and substitutions. The idea is to try to find the set of operations that can transform our average image into a given input image. Then instead of storing the input image itself, we could store only the sequence of operations that can reproduce it from the average image.

While studying the theory of minimum string edit distance, a new research direction was opened: “image morphing”. Image morphing refers to the transformation of one image into another. The idea is to use image morphing techniques in order to transform a statistical image (e.g. the “average” image) into a given input image. Then we could store only the transformation parameters instead of the input image. The more similar the images are, the simpler the transformation would be, therefore the less storage space it would require. This is a very promising research direction for SRC compression, and future research efforts should explore it thoroughly.

9. CONCLUSION

This research project produced three new SRC compression methods. These methods can be used to compress sets of similar images delivering performance that is comparable to well known state-of-the-art compression methods.

An advantage of SRC methods over other compression methods is that they offer increased security: when images are SRC-compressed, they are also automatically encrypted, by having as “key” the statistical image used (i.e. the median or max image). This key image can be stored separately. Unlike other compression methods where a compressed image can always be decompressed by a third party with a decompressor, with SRC methods someone needs not only the decompressor, but also the key statistical

image. Therefore SRC-compressed images can be securely transferred through computer networks, since the key image is not part of the transmission (it is stored permanently at the origin and destination).

Further research should explore the use of SRC methods as combined compression+encryption schemes. Also, as it has been mentioned above, further research should try to incorporate block-matching into the SRC methods, and also develop new or improve current SRC methods using the ideas of “minimum string edit distance” and image morphing.